arXiv:1208.5246v1 [physics.comp-ph] 26 Aug 2012

PARAMETERIZATION OF NON-LINEAR MANIFOLDS

C. W. GEAR
DEPARTMENT OF CHEMICAL AND BIOLOGICAL ENGINEERING
PRINCETON UNIVERSITY, PRINCETON, NJ
E-MAIL:WGEAR@PRINCETON.EDU

Abstract. In this report we consider the parameterization of low-dimensional manifolds that are specified (approximately) by a set of points very close to the manifold in the original high-dimensional space. Our objective is to obtain a parameterization that is (1-1) and non singular (in the sense that the Jacobian of the map between the manifold and the parameter space is bounded and non singular).

Keywords Diffusion Maps, PCA, Distance Matrices

1. Introduction. When it is known that the high-dimensional solutions of a computational problem lie on a lower-dimensional manifold we may wish to operate only on that manifold. By repeatedly solving the problem numerically, we can compute points that are approximately on the manifold and we will assume that we are able to compute a set of such points over the manifold so that they are "reasonably distributed" (i.e., the average density of the points does not vary too greatly from region to region. We also assume that we can approximate any point on the manifold sufficiently accurately by interpolation from these pre-computed points and we would like to find a convenient parameterization of the manifold that can be used for the interpolation.

For example, if the solutions of a differential equation in the high-dimensional space rapidly approach a low-dimensional manifold it may be desirable to integrate a reduced system on the manifold to avoid stiffness. In this case, it might be even more convenient to integrate in the parameter space. Note that in this example it is important that the parameterization be non singular: otherwise the differential equations in the parameter space would have a singularity.

We suppose that we are given a set of N points X_i, $i = 1, \cdots, N$ in s-dimensional Euclidean space, S, with coordinates $\mathrm{x}_i = \{x_{q,i}\}$, $q = 1, \cdots, s$ that are on (or close to, due to noise or computational error) a low-dimensional manifold, M. Our ultimate objective is to find, computationally, a parameterization of any manifold that can be simply parameterized. For now we rule out manifolds such as tori that can only be parameterized by using periodicity in some of the parameters. A desiderata is to be able to parameterize any manifold that can be formed by "twisting" (without any stretching) a linear manifold in the high dimensional space (such as the well-known "Swiss roll" example in Figure 1.1(b)) so that the parameterization is a good approximation to a Cartesian coordinate system for the original linear manifold. (In that figure the X-axis is foreshortened and has a range of about 2.5 times that of the other two axes so as to show the structure. Consequently it appears to be rolled from a fairly narrow rectangle but it is actually rolled from the square shown in Figure 1.1(a). The randomness of these points was constrained to get a "reasonable" distribution - the 1 by 1 square was divided into 40 by 40 equal smaller squares, a point was placed in the middle of each and then subject to a random uniform perturbation in each coordinate in the range [-0.4, 0.4]/40, thus ensuring that no points were closer than 0.2/40.)

The learning and statistical communities have long been interested in manifold learning and representation so there is a large body of literature. In those cases the typical task is to start with the data set $\{\mathrm{X}_i\}$ and produce a "representation" in a lower-dimensional space such that the differences in the distances of pairs of points in the two spaces are suitably small. This was achieved by minimizing a *stress* function in Sammon's early paper[3] and numerous modifications have been proposed, such as [5, 6]. If the low-dimensional

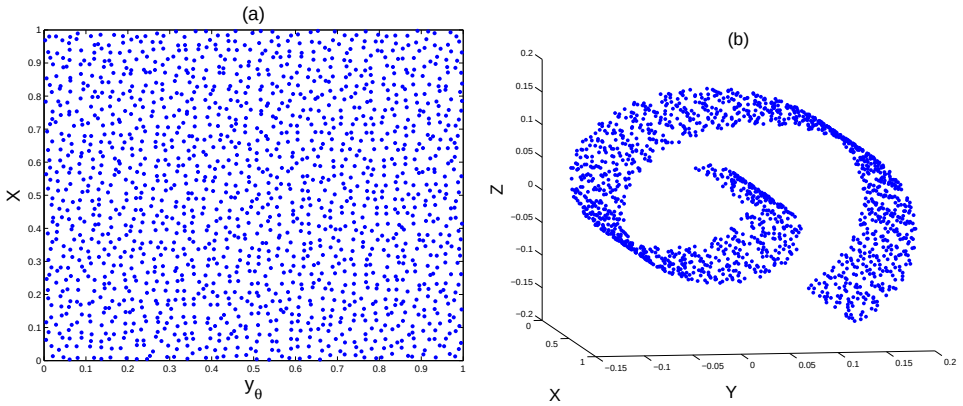

FIG. 1.1. *(a) 1,600 points "randomly" placed (see text) on a 1 by 1 square. (b) "Swiss Roll" formed by rolling the square into a spiral. The y_θ coordinate in (a) correspons to distance around the spiral in (b).*

manifold is linear, the Principal Component Analysis (PCA) is the standard way to identify and represent it - the principal components provide an orthogonal basis for the smallest linear space containing the manifold. Many methods for non-linear manifolds start with some modification of PCA. For example, [8] effectively uses PCA locally and then stitches the tangent spaces together. Local Linear Embedding (LLE) [2] looks for a local linear approximation which is related to PCA.

Another class of methods is based on the structure of the graph generated by focussing attention primarily on nearby neighbors. Diffusion Maps [1] are one type of this method. We have looked at the use of Diffusion Maps for the parameterization of slow manifolds [4]. Diffusion Maps use a parameter, ϵ, that effectively specifies the distance of what is considered a nearby neighbor. When a small ϵ is used, the parameterization is nearly singular at the boundary. A much larger ϵ sometimes overcomes this difficulty, but can lead to other difficulties.

In this report we show that that the limit of a particular case of diffusion maps for large ϵ is PCA if we use the left eigenvectors rather than the right eigenvectors normally used in diffusion maps. However, because, in diffusion maps, we work with distances in the graph structure rather than with the coordinates $\{x_i\}$ that are the input to PCA, we can consider modifications to the distances of pairs of points that are not nearby neighbors to improve the parameterization.

In Section 2 we discuss the problems of diffusion maps for the data set in Figure 1.1(b) - for small ϵ the map has a nearly singular Jacobians, but breaks down for large ϵ because points distant in the manifold are relatively close in the high-dimensional space. Then, in Section 3, we consider the large ϵ limit for diffusion maps on linear manifolds and prove that the left eigenvectors provide the PCA decomposition and show that this can be obtained by working directly with the distance matrix. In Section 4 we consider ways to modify the distance matrix to improve the parameterization.

2. Diffusion Map Issues. Define the (symmetric) distance matrix $\mathbf{H}$ to have the entries

$$\mathbf{H}_{ij} = d(X_i, X_j)$$

where $d(.,.)$ is the Euclidean distance between the points, that is

$$\mathbf{H}_{ij} = \sqrt{\sum_{q=1}^{s}(x_{q,i} - x_{q,j})^2}.$$

The diffusion map method [1] now forms the matrix $\mathbf{W} = \{\mathbf{W}_{ij}\} = \{w(\mathbf{H}_{ij})\}$ where w is a suitable non-negative function such that $w(0) = 1$ and $w(h)$ is monotonic non-increasing for $h > 0$ with $w(\infty) = 0$. We will use

$$w(h) = \exp(-\left(\frac{h}{\epsilon}\right)^2).$$

The diffusion map method continues by dividing each row of $\mathbf{W}$ by its corresponding row sum to get a Markov matrix $\mathbf{K}$ whose largest eigenvalue is 1 with a corresponding eigenvector $\mathbf{e} = [1, 1, \cdots, 1]^T$. Some of its next few eigenvectors serve to identify the low-dimensional manifold and the i-th components of these eigenvectors provide a parameterization for the points, X_i, on the manifold. It is well known that for small ϵ the parameterization is poor near the boundaries - see, for example, [4] - but that a large ϵ can give a good parameterization under some circumstances.

In Figure 2.1 we show the results of applying a diffusion map to the data in Figure 1.1 (b). Here, ϵ was chosen as $\sqrt{2}\beta$ where β was the smallest value such that if all edges in the graph corresponding to the distance matrix $\mathbf{H}$ longer than β were removed, the graph would still be connected. (For this particular data set, β is 0.0322.) Figure 2.1 (a) and (b) plot the components of the second and third eigenvectors, ξ_2 and ξ_3, against the X and y_θ coordinates of the points in Figure 1.1(b). (X and y_θ are the coordinates of the points in Figure 1.1(a). After the square is "rolled" y_θ is the arc length around the spiral from its inner stating point while X is the coordinate along the length of the cylindrical spiral.) We see that ξ_2 and ξ_3 come very close to providing direct parameterizations of X and y_θ, respectively, and this is shown more clearly in Figure 2.1 (c) and (d) (which are projections of the previous two figures onto axial panes). However, we see that the parameterization is very poor near the boundary where $\partial\xi_2/\partial X$ and $\partial\xi_3/\partial y_\theta$ are nearly zero (because the curve is close to a cosine half cycle).

In some cases, this can be overcome by using a large ϵ. In Figure 2.2 (a) and (b) we show the equivalent of Figure 2.1 (c) and (d) for a diffusion map of the same data with $\epsilon = 1$ - the distance along the cylinder. Note that ξ_2 provides a very good parameterization of the X direction - it is almost linear, but that ξ_3 is no longer a (1-1) map. This occurs because $\mathbf{W}_{ij}$ for points on adjacent turns of the spiral is about 0.985 so the points are seen as close. (The spiral started with radius 0.03 and increased radially by 0.12 on each full turn.)

We would like to get the benefits of a large ϵ - the linearity in Figure 2.2(a) without the drawbacks illustrated in Figure 2.2(b). For that reason, we study the large ϵ limit of diffusion maps in the next section.

3. Diffusion Maps in the Large ϵ Limit. We are really interested in non-linear manifolds, but for expository purposes we start by assuming that the manifold is linear. Set the (arbitrary) origin to be at the point mean. This implies that

$$\sum_{i=1}^{n} x_{q,i} = 0 \tag{3.1}$$

for all q. Since the manifold is linear, it is now a linear subspace, D.

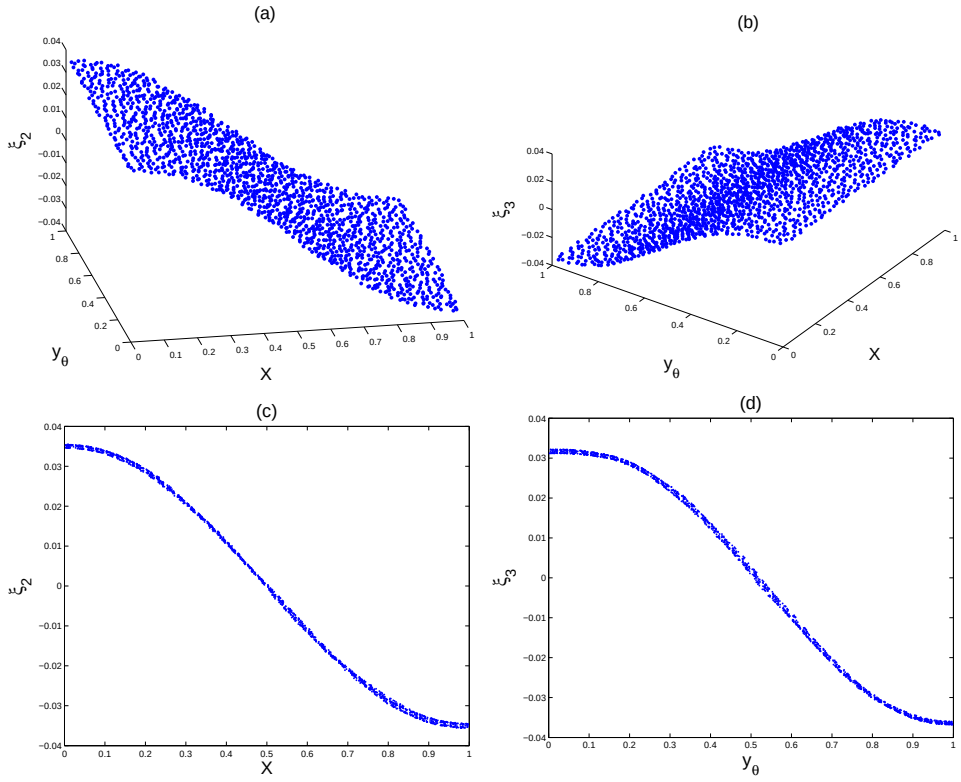

FIG. 2.1. *Diffusion map applied to data in Figure 1.1(b) with* $\epsilon = 0.0455$. *(a) Second eigenvector* ξ_2. *(b) Third eigenvector* ξ_3. *(c)* ξ_2 *versus* X. *(d)* ξ_3 *versus* y_θ.

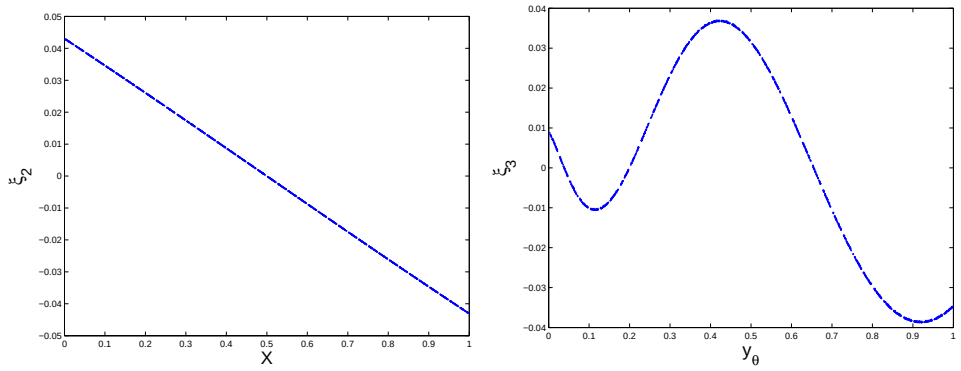

FIG. 2.2. *Diffusion map applied to data in Figure 1.1(b) with* $\epsilon = 1$. *(a)* ξ_2 *versus* X. *(b)* ξ_3 *versus* y_θ.

Define $\mathbf{x}_q$ to be the vector whose i-th component is $x_{q,i}$ and $\mathbf{e} = [1, 1, \cdots, 1]^T$ so eq (3.1) can be written as

$$\mathbf{e}^T \mathbf{x}_q = 0 \tag{3.2}$$

for all q. (Bold face lower-case roman letters will always refer to N-dimensional column vectors while non-bold face roman letters will refer to s-dimensional column vectors.)

Suppose ϵ is large enough that $(\max(\mathbf{H}_{ij})/\epsilon)^4$ can be ignored. In that case,

$$\mathbf{W}_{ij} \approx 1 - \left(\frac{\mathbf{H}_{ij}}{\epsilon}\right)^2$$

The row sums of this matrix are

$$\sum_{j=1}^{N} \mathbf{W}_{ij} \approx N - \frac{1}{\epsilon^2} \sum_{j=1}^{N} \mathbf{H}_{ij}^2$$

so

$$\mathbf{K}_{ij} \approx \frac{1}{N}(1 - \frac{1}{\epsilon^2}[\mathbf{H}_{ij}^2 - \frac{1}{N}\sum_{k=1}^{N} \mathbf{H}_{ik}^2]) \tag{3.3}$$

In other words,

$$\mathbf{K} \approx \frac{1}{N}(\mathbf{E} - \mathbf{F})$$

where $\mathbf{E}$ is the matrix of all ones and

$$\begin{aligned}
\mathbf{F}_{ij} &= \frac{1}{\epsilon^2}[\mathbf{H}_{ij}^2 - \frac{1}{N}\sum_{k=1}^{N} \mathbf{H}_{ik}^2]) \\
&= \sum_{q=1}^{s}[(x_{q,i} - x_{q,j})^2 - \sum_{k=1}^{N}(x_{q,1} - x_{q,k})^2/N] \\
&= \sum_{q=1}^{s}[x_{q,j}^2 - 2x_{q,i}x_{q,j} - \sum_{k=1}^{N} x_{q,k}^2/N].
\end{aligned} \tag{3.4}$$

Writing c_q for $\sum_{k=1}^{N} x_{q,k}^2$ and defining vector $\mathbf{v}_q$ to have i-th entries $x_{q,i}^2$, eq (3.4) can be written as

$$\mathbf{F} = \sum_{q=1}^{s}[\frac{\mathbf{e}(\mathbf{v}_q - c_q\mathbf{e})^T}{N} - 2\mathbf{x}_q\mathbf{x}_q^T]. \tag{3.5}$$

The points X_i lie in the linear subspace D. Chose a coordinate system for S such that the first d coordinates span D. Hence, for $q > d$ $\mathbf{x}_q = \mathbf{v}_q = 0$ and $c_q = 0$ so we can rewrite eq (3.5) as

$$\mathbf{F} = \sum_{q=1}^{d}[\frac{\mathbf{e}(\mathbf{v}_q - c_q\mathbf{e})^T}{N} - 2\mathbf{x}_q\mathbf{x}_q^T]. \tag{3.6}$$

Note that if a left or right eigenvector, $\mathbf{v}$ of $\mathbf{K}$ given by eq (3.3) satisfies $\mathbf{e}^T\mathbf{v} = 0$ then it is also an eigenvector of $\mathbf{F}$and *vice versa*.

Observation

If the points $\{X_i\}$ are regularly placed in the linear subspace - that is is they lie on the points of a regular rectangular grid or its extension to higher dimensions - and if the coordinate system is aligned with this grid, then $\mathbf{x}_q, q = 1, \cdots, d$ are eigenvectors of $\mathbf{F}$ and hence $\mathbf{K}$.

This follows by construction. In this case the vector $\mathbf{x}_q$ takes the form

$$\mathbf{x}_q = e_1 \otimes e_2 \otimes \cdots \otimes e_{q-1} \otimes z_q \otimes e_{q+1} \otimes \cdots \otimes e_d \tag{3.7}$$

where e_i is a column vector of ones of length equal to the number of grid points in the i-th dimension and z_q is the set of coordinate points in the q-th dimension. Note that $e_q^T z_q = 0$ for all q so that the $\mathbf{x}_q$ are mutually orthogonal and all orthogonal to $\mathbf{e}$. Hence

$$\begin{aligned}\mathbf{F}\mathbf{x}_q &= \sum_{p=1}^{d} [\frac{\mathbf{e}(\mathbf{v}_p - c_p\mathbf{e})^T}{N} - 2\mathbf{x}_p\mathbf{x}_p^T]\mathbf{x}_q \\ &= \sum_{p=1}^{d} \frac{\mathbf{e}\mathbf{v}_p^T}{N}\mathbf{x}_q - 2(\mathbf{x}_q^T\mathbf{x}_q)\mathbf{x}_q. \end{aligned} \tag{3.8}$$

If $p \neq q$ then $\mathbf{v}_p^T\mathbf{x}_q = 0$ directly from eq (3.7) so it only remains to show that for this data,

$$\mathbf{v}_q^T\mathbf{x}_q = \sum_{i=1}^{N} x_{q,i}^3 = 0$$

which is true by virtue of the regular spacing and $\mathbf{e}^T\mathbf{x}_q = 0$.

If the points are "reasonably" distributed, then this property is approximately true. In Figure 3.1 we show the points plotted via their entries, $\xi_{2,i}$ and $\xi_{3,i}$ of the eigenvectors of $\mathbf{K}$ for the planar data in Figure 1.1(a). Because the eigenvalue computation finds the dominant axis as second leading eigenvector (after the eigenvector $\mathbf{e}$) and that tends towards one of the diagonals when the data is random, the square has been rotated. To estimate how far this is from being an exact parameterization of the data, we computed the best (in the least squares sense) affine transformation from the (ξ_2, ξ_3)-plane to the (X, y_θ)-plane of Figure 1.1(a) and computed the maximum L_2 error in the (X, y_θ)-plane over all points. The result was 1.7899×10^{-14} - only barely larger than round-off errors.

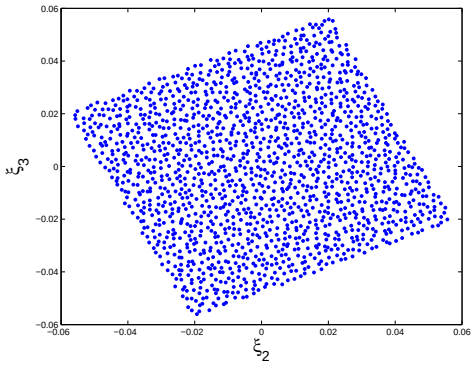

FIG. 3.1. *Points plotted by right eigenvectors of* $\mathbf{F}$.

However, it is more important to note that the left eigenvectors of $\mathbf{F}$ always have this property, regardless of the distribution of points, so we propose to use the left eigensystem of $\mathbf{F}$ but to avoid having to repeat "left eigenvector" we will instead work with $\mathbf{G} = \mathbf{F}^T$ whose eigenvectors $\mathbf{u}$ are the left eigenvectors of $\mathbf{F}$.

We are dealing with vectors in two different spaces: the s-dimensional original Euclidean space containing the points whose coordinates are $\mathrm{x}_i = [x_{1,i}, x_{2,i}, \cdots, x_{s,i}]^T$ and the N-dimensional space, P, containing the vectors $\mathbf{x}_q = [x_{q,1}, x_{q,2}, \cdots, x_{q,N}]^T$. If we pick any basis for the Euclidean space (which has to be orthogonal so that eq (3.4) is based on the distance) it determines the entries in the N-dimensional vectors $\mathbf{x}_q$. For each orthogonal set of vectors in the Euclidean space (which we will call *geometric* vectors) there is a corresponding set of vectors of the point coordinates (which we will call *point* vectors). In particular, if we pick an (orthogonal) basis set for S such that the first d basis vectors span D, then the $x_{q,i}$ satisfy $x_{q,i} = 0, q > d$. Hence, the remaining $s - d$ basis vectors do not affect the d vectors $\mathbf{x}_q, 1 \leq q \leq d$. In this sense, to every orthogonal basis for D there is a corresponding set of d linearly-independent point vectors in P.

To avoid verbal complexity, we will refer to "non-zero eigenvectors" to mean "eigenvectors corresponding to non-zero eigenvalues" and its obvious extensions.

We find that G has the interesting properties summarized in

Theorem

If the components of $\mathbf{G}$ *are given by*

$$\mathbf{G} = \sum_{q=1}^{d} [\frac{(\mathbf{v}_q - c_q \mathbf{e})\mathbf{e}^T}{N} - 2\mathbf{x}_q \mathbf{x}_q^T] \tag{3.9}$$

then

1. $\mathbf{G}$ has exactly d non-zero eigenvalues.
2. The zero eigenvalue has algebraic multiplicity $N - d$ and geometric multiplicity at least $N - d - 1$ although there are special cases where the geometric multiplicity is also $N - d$.
3. There exists a geometrically orthogonal basis for D such that the d corresponding non-null point vectors $\mathbf{x}_q$ are the non-zero eigenvectors of $\mathbf{G}$.
4. These eigenvectors are orthogonal.
5. The corresponding eigenvalues are $-2\mathbf{x}_q^T \mathbf{x}_q$ so that all non-zero eigenvalues are negative.

Proof

We restrict ourselves to coordinate systems in S such that the first d cordinates span M. From eq (3.9) we see that if $\mathbf{b}$ is any vector orthogonal to $\mathbf{e}$ and $\mathbf{x}_q,\ q = 1, \cdots, d$ then $\mathbf{G}\mathbf{b} = 0$ so $\mathbf{b}$ is a zero eigenvector. Hence any $s - d - 1$ (linearly independent) vectors that span the space orthogonal to $\mathbf{e}$ and $\mathbf{x}_q,\ q = 1, \cdots, d$ are independent zero eigenvectors and are orthogonal to $\mathbf{e}$. Note that $\mathbf{e}^T$ is a zero left eigenvector of $\mathbf{K}$. Since it is orthogonal to the $s-d-1$ zero right eigenvectors just enumerated, there must be one more zero eigenvalue. (Usually its right eigenvector is a generalized eigenvector but there are special cases when this is not so, for example, when $N - 1 = d = s$.)

We complete the proof by showing that there are d independent non-zero eigenvectors.

Let $\mathbf{X}$ be the matrix whose columns are $\mathbf{x}_q, q = 1, \cdots, d$, and let $\mathbf{y} = \mathbf{X}\alpha$ be a linear combination of those vectors where α is a d-dimensional vector. Hence from eq (3.9) and $\mathbf{e}^T X = 0$ we have

$$\mathbf{G}\mathbf{y} = -2\mathbf{X}\mathbf{X}^T\mathbf{X}\alpha \tag{3.10}$$

Define

$$Y = \mathbf{X}\mathbf{X}^T. \tag{3.11}$$

It is a d by d symmetric matrix so has real eigenvalues and mutually orthogonal eigenvectors. Because the points are in the d-dimensional subspace D and no lower dimensional subspace, Y is positive definite so all of its eigenvalues are strictly positive. Let α be one of the eigenvectors with corresponding eigenvalue μ. Hence

$$\mathbf{G}\mathbf{y} = -2\mu\mathbf{X}\alpha = -2\mu\mathbf{y}$$

so $\mathbf{y} = \mathbf{X}\alpha$ is an eigenvector of $\mathbf{G}$ with eigenvalue -2μ. This provides a complete set of d non-zero eigenvectors. If Q is an orthonormal matrix whose columns are the (scaled) eigenvectors, $\{\alpha\}$, of Y then Q can be applied to the initially selected Cartesian coordinate system for D to get a new Cartesian coordinate system with property 3.

Property 4 follows because the action of $\mathbf{G}$ defined by eq (3.9) on the subspace spanned by $\{\mathbf{x}_q\}$ is exactly the same as the action of

$$\hat{\mathbf{G}} = -2\sum_{q=1}^{d} \mathbf{x}_q \mathbf{x}_q^T \tag{3.12}$$

$\hat{\mathbf{G}}$ is a symmetric matrix so its eigenvectors are orthogonal. Property 5 follows from eq (3.9) when we choose a basis for D such that the eigenvectors are $\{\mathbf{x}_q\}$ (we have demonstrated how to do above). Then Y is diagonal with entries $\mathbf{x}_q^T\mathbf{x}_q$ which are its d eigenvalues μ_q.

As a footnote to this proof we observe that the number of non-zero eigenvalues follows quickly from a paper of nearly 75 years ago by Young and Householder[7]. They show that the matrix

$$\begin{bmatrix} \mathbf{H} & \mathbf{e} \\ \mathbf{e}^T & 0 \end{bmatrix}$$

has rank $d+2$ (see eq. (2) in their paper). If we subtract the average of the first N rows from the last and then add the new last row to each of the first N rows (not changing the matrix rank) we get the matrix

$$\begin{bmatrix} \mathbf{E}+\mathbf{G} & \mathbf{0} \\ \mathbf{z}^T & -1 \end{bmatrix}$$

where $\mathbf{z}$ is some N-dimensional vector whose value does not affect the rank. Hence, $\mathbf{E}+\mathbf{G}$ has rank $d+1$. Since $\mathbf{e}^T$ is a left eigenvector of $\mathbf{G}$ and $\mathbf{E}$ with eigenvalues 0 and 1 respectively and $\mathbf{E}$ has rank 1, the rank of $\mathbf{G}$ is d.

4. Modifying Distances to get a Good Parameterization. We see from eq (3.9) that the non-zero eigenvectors of $\mathbf{G}$ are the eigenvectors of XX^T and these are precisely the principal components from PCA so the method proposed in Section 3 provides exactly the PCA representation of the data set, which raises the question "What is the advantage of this method over PCA?" PCA starts with the coordinates of the data whereas this method starts with the distances of each point pair (and incidentally provides the coordinates in an orthogonal coordinate system from the calculation). The problems with the diffusion map parameterization (using large ϵ as discussed in Section 2) arose when distances that were large in the manifold were small in the original space. Because the algorithm only uses distances we can consider modifying the distances of more distant objects to be approximately compatible with those in the manifold.

A procedure for doing this is as follows:

1. Remove all edges from the graph representation of the distance matrix $\mathbf{H}$ longer than some value.

2. Compute the new shortest path to all edges to get a modified distance matrix
3. Use this modified matrix to generate a parameterization using the technique of Section 3

In Figure 4.1(a) we show the points plotted in the (ξ_1, ξ_2)-plane using the modified distance matrix approach for the Swiss roll data in Figure 1.1(b). Edges were removed from the graph that were longer than 2β where β was the minimum value that did not lead to a disconnected graph when all larger edges were cut (β is 0.0322). As happened in Figure 3.1, the square is twisted. As before, we found the best affine transformation of these points to match the data in Figure 1.1(a). The new coordinate systems is called ϕ_1 and ϕ_2. In Figure 4.1 (b) and (c) we show the plots of X versus ϕ_1 and y_θ versus ϕ_2. As can be seen, the relationship is almost linear. The largest distance between the rotated data and the data on the original square before it was "Swiss rolled" was 1.7899×10^{-14}, again, barely larger than round off.

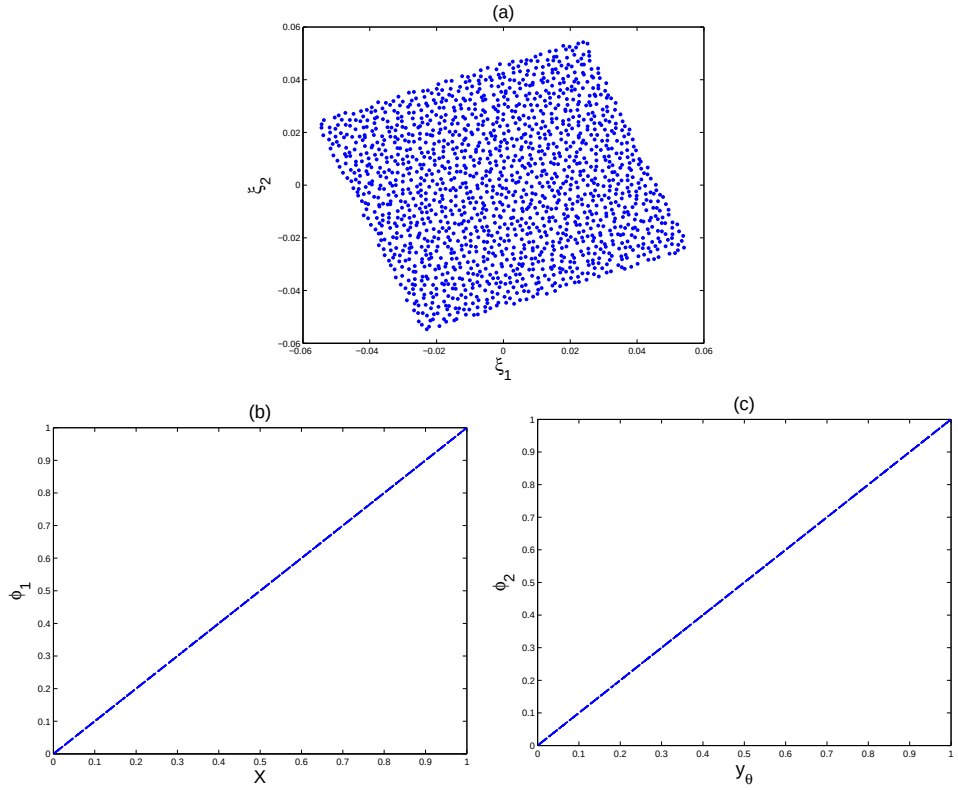

FIG. 4.1. *Modified distance methods applied to Swiss roll data. (a) Points plotted by right eigenvectors of* $\mathbf{G}$. *(b) and (c) Plot of rotated coordinates of (a) against coordinates in original plane data.*

The distance modification scheme gives good results for the Swiss roll data because it was obtained by twisting a linear manifold without stretching. However, if it is necessary to stretch the data to "unfold" it into a linear manifold, the scheme will not work unless some of the nearest neighbor distances are also modified. This can be seen with the example in Figure 4.2(a) which consists of a set of points somewhat uniformly placed on the majority of a surface of a sphere. In this example, the surface of the sphere is included from the "North pole" to to Southern latitude 67.5^o. A point was placed at the North pole and then

points were place on each of 15 circles of constant latitude at constant separations of 10.5^o. The points were equi-spaced on each circle. 40 points were placed on the equator and the number on each of the other circles was chosen to make the spacing on each circle as similar as possible except that a minimum of 6 points were placed on each circle. The location of one of the points was chosen from a uniform random distribution. In Figure 4.2(a) all points are plotted as dots except for the points on the southern-most latitude which are circles. Clearly this 2D surface can be parameterized easily. The southern-most latitude circle is the boundary of the finite, non-linear manifold.

The results of the standard diffusion map and the modified distance scheme are shown in Figure 4.2 (b) and (c), respectively. The plots shown the circles of constant latitude in the eigenvalue coordinates. The southern-most circle is shown as a thicker line.

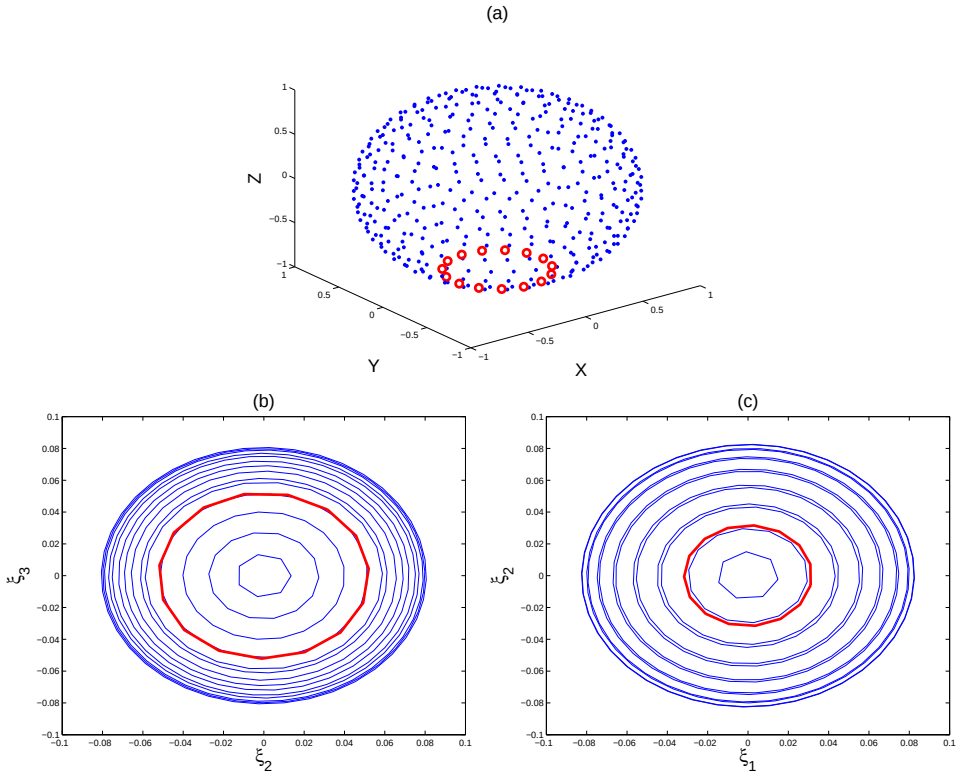

FIG. 4.2. *(a) Original data on most of the surface of a sphere. (b) Diffusion Map applied to this data. (c) Modified distance scheme applied to same data.*

As can be seen from the figure, the modified distance scheme has the boundary of the manifold in the interior of parameterization it produces (as does the diffusion map process, although not as badly). This arise because the scheme places a great emphasis preserving the length of nearby neighbors, and in this case there is a circle of nearby neighbors with a perimeter that is much smaller than the perimeter of other circles that are in the interior of the 2D manifold. Clearly it is necessary to modify the lengths of some nearby neighbors to get a good parameterization. A clue to the need for this can be obtained from the size of additional eigenvalues of $\mathbf{G}$ and local changes in the components of their eigenvectors. Since

the method is equivalent to a PCA decomposition[1], other eigenvalues must be large when we find that points significantly separated in the original space are close in the parameter space. How to systematically modify nearby neighbors distances in this case is a subject of ongoing research.

REFERENCES

[1] R. R. Coifman and S. Lafon, *Appl. Comput. Harmon. Anal.* **21**, pp. 5–30 (2006)

[2] S. T. Roweis and L. K. Saul, Nonlinear Dimensionality Reduction by Locally Linear Embedding, *Science*, **290**, pp 2323-2326, (2000)

[3] J. W. Sammon Jr., A Nonlinear Mapping for Data Structure Analysis, *IEEE Trans. Computers*, **C-18**, 5, pp 401-409, (1969)

[4] B. Sonday, C. W. Gear, A. Singer and I. G. Kevrekidis, Solving Differential Equations by Model Reduction on Learned Manifolds, see www.princeton.edu/∼wgear/SGSK.pdf

[5] J. Sun, M. Crowe, and C. Fyfe, Extending Metric Multidimensional Scaling with Bregman Divergences, *Pattern Recognition*, **44**, pp 1137-1154, (2011)

[6] J. Sun, C. Fyfe, and M. Crowe, Extending Sammon Mapping with Bregman Divergences, *Information Sciences*, **187**, pp 72-92, (2012)

[7] G. Young and A. S. Householder, Discussion of a Set of Points in Terms of their Mutual Distances, *Psychometrica*, **3**, 1, pp 19-21, (1938)

[8] Z. Zhang and H. Zha, Principal Manifolds and Nonlinear Dimension Reduction via Local Tangent Space Alignment, *SIAM Journ. on Scientific Computing*, **268**, 1, pp 313-338, (2005)

[1] This is true if the distance matrix corresponds to a feasible configuration in real space. If it doesn't, for example because the distance inequality is not satisfied, we will get some complex eigenvalues and eigenvectors